ABSTRACT OF THE DISCLOSURE

A parallel hash ripple join algorithm partitions tuples of two relations for localized processing. The algorithm is non-blocking and may be performed in a parallel, multi-processor environment. At each processing node, the tuples are further partitioned such that join operations may be performed as tuples are redistributed to each node during the partitioning.